BeSpaceD

Towards a Tool Framework and Methodology for the Specification and Verification of Spatial Behavior of Distributed Software Component Systems

Jan Olaf Blech and Heinz Schmidt

RMIT University, Melbourne, Australia

arXiv:1404.3537v1 [cs.SE] 14 Apr 2014

In this report, we present work towards a framework for modeling and checking behavior of spatially distributed component systems. Design goals of our framework are the ability to model spatial behavior in a component oriented, simple and intuitive way, the possibility to automatically analyse and verify systems and integration possibilities with other modeling and verification tools. We present examples and the verification steps necessary to prove properties such as range coverage or the absence of collisions between components and technical details.

1 Introduction

Ensuring that distributed embedded systems meet their requirements is an important and well-studied topic in industry and research. With growing system sizes and the combination of independently developed systems into larger systems of systems the interplay of different subsystems becomes an important aspect of study. When regarding embedded controllers integrated into other devices such as robots or cars, the software running on these systems does have an impact on the physical environment. Potential hazards resulting from incorrect software can lead to damaged property, injuries or even loss of human lives. When combining systems that have a physical impact on an environment, one may have to study the subsystems together, since one can detect some potential hazards, e.g., two cars driving in opposite directions, but on a single lane only in the combined system.

Different techniques have been introduced for modeling and checking of cyber-physical systems and their properties. These comprise differential equations, automata and different notions of time such as continuos or discrete time and others. Based on these modeling approaches, tools for checking properties have been developed. The software engineering community has been studying the design and architecture of component-based systems for decades. Specification formalisms like automata and message sequence charts are frequently used to describe expected behavior of systems. Different means for checking these specifications are available.

In this report, we present work towards a unified approach to cyber-physical systems and component-based software development. We motivate a new framework for spatial cyber-physical behavior, communication and interaction between different components. Our framework is especially suited for large-scale widely distributed systems with limited interactions between components but also allows the modeling of the precise geometric behavior of components in space. Several components may be distributed over a large space and have distinct features of communication, spatial and internal behavioral aspects. We describe a process for checking properties of our models. In particular, we are focusing on spatial aspects: a component features different dimensions of possible non-deterministically occupied space, e.g., characterizing its physical dimensions, and the reach of sensors and communication devices. We study modeling and verification scenarios which are characterized by a discrete notion of time that features *time points* and *time intervals* between them. Time points are partially ordered. This allows for the modelling of synchronous and asynchronous systems with distinct synchronisation points between different components.

Some preliminary ideas to this report are presented in [9]. An application of our framework in an IDE for the development of reactive probabilistic systems is featured among other content in [20].

Overview Section 2 discusses related work. Modeling of systems with our framework and examples are presented in Section 3. Section 4 describes ways to reason about our models and perform verification. Section 5 discusses our implementation in Scala. Conclusion and future work is presented in Section 6.

2 Related Work

Work that is relevant to this paper has been done in areas such as formal logic and process algebras, hybrid-systems, robotics and formal methods for component-based software engineering.

Formal Logic and Process Algebraic Approaches The handbook of spatial logic [1] discusses spatial logics, related algebras and applications. The book covers a large spectrum not only limited to computer science. A process algebra like formalism for describing and reasoning about spatial behavior has been introduced in [11] and [12]. Process algebras come with a precise formal semantics definition and are aimed towards the specification of highly parallel systems. Here, disjoint logical spaces are represented in terms of expressions by bracketing structures and carry or exchange concurrent process representations. A model checking tool for process algebra like spatial behavioral specifications is presented in [13]. A graph-based technique for the verification of spatial properties of finite π-calculus fragments is introduced in [18].

For results on spatial interpretations see, e.g., [21]. Many aspects of spatial logic are in general undecidable. A quantifier-free rational fragment of ambient logic (corresponding to regular language constraints), however, has been shown to be decidable in [36].

Special modal logics for spatio-temporal reasoning go back to the seventies. The Region Connection Calculus (RCC) [4] includes spatial predicates of separation. For example RCC features predicates indicating that regions do not share points at all, points on the boundary of regions are shared, internal contact where one region is included and touches on the boundary of another from the inside, proper overlap of regions, and proper inclusion. In addition [4] features an overview of the relation of these logics to various Kripke-style modal logics, reductions of RCC-style fragments to a minimal number of topological predicates, their relationship to interval-temporal logics and decidability.

Hybrid-systems The area of hybrid systems has seen the development of different tools for reasoning and verification. SpaceEx [28] allows the modeling of continuos hybrid systems based on hybrid automata. It can be used for computing overapproximations of the space occupied by an object moving in time and space. Additionally, it is possible to model spatial behavior in more general purpose oriented verification tools in Hybrid systems (e.g., [31]).

Robotics Related to our work is the work on path planning for robots (e.g., [23, 25]). In our work, however, we are concentrating on checking existing properties of systems rather than optimization or discovery of new possible paths. Collision detection for robots in combination with motion planning has been studied for a long time, see, e.g., [22] and [14]. Strongly related to motion planning is the task of efficiently handling geometric reasoning. On this geometric interpretation level, techniques have been investigated to structure the tasks of detecting possible inference between geometric objects (e.g., [19] and [26]) for efficient analysis.

Data Models and GIS like Services Data models for cyberphysical infrastructure in construction, plant automation and transport – domains that we are aiming for in this paper – have been studied in the past. Unlike in this paper, many of the existing real-world applications are aligned towards a geometric representation of components and are typically based on so-called 2.5 dimensional GIS (Geographic Information System) representations where the 3rd dimension $z = f(x, y)$ is represented as a function f of the 2 dimension x and y coordinates. This [2] limits the geometric, topological and information retrieval use of such models. True three dimensional modeling is far from common practice [33]. Our approach is not limited to a particular geometric representation, coordinate or dimension system. Future extensions of interest include consideration of standards such as the Web 3D Services and the Sensor Web Enablement Architecture of the Open Geospatial Consortium [1], visualisation and decision support [35] and efficient data structures for fast meta reasoning and presenting subproblems of choice to specialist solvers used in our examples.

Component-based Systems Multiple formalisms to describe component behavior can be used as specification basis for our methodology. For example the UML 2.0 [10]

[1] `http://www.opengeospatial.org`

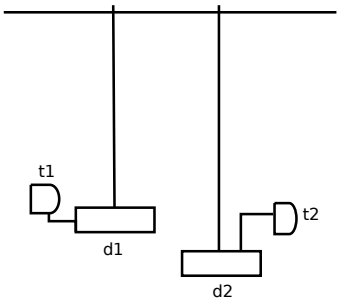

Figure 1: Concurrent window cleaning

standard provides message sequence charts and state machines that can be used to describe component behavior and interactions. These can be used to derive information about shared lifetimes of components and behavior.

The notion of time and component interaction used in this paper is compatible with the Petri-nets induced notion. This is used, e.g., in the BIP framework [3] for modeling of distributed asynchronous systems. Invariants are used as an intermediate step for verification in the BIP context [5]. Invariants are also used as an intermediate representation of components in this work.

We have been investigating mathematical models of behavior in space in previous work [34], which also features a survey on work related to these models. A prestudy of the work described in this paper is presented in [9]. Another software architectural model was proposed in [32]. Specification and reasoning on existing software component systems [7, 6] in the context of a type system for behavior has also been examined by us and is an important direction of future extensions of this paper. In this work we aim towards a unified view of component system aspects and verification and the introduction of space as an additional feature. Symbolically reasoning about invariants of asynchronous distributed systems, which is a part of our verification methodology in this paper, has been studied by us in [8].

3 Modeling of Spatial Behavior

This section presents the ingredients of our modeling framework for spatial behavior and examples.

3.1 Guiding Examples

Here, we present four guiding examples, concurrent window cleaning, moving forklifts, communicating cars on a road network, and rotating robots.

Concurrent Window Cleaning Figure 1 shows a concurrent window cleaning system. Platforms d_1, d_2 are attached to a mobile device on a roof. They can be moved horizontally by moving the mobile device on the roof and vertically through a rope. Attached

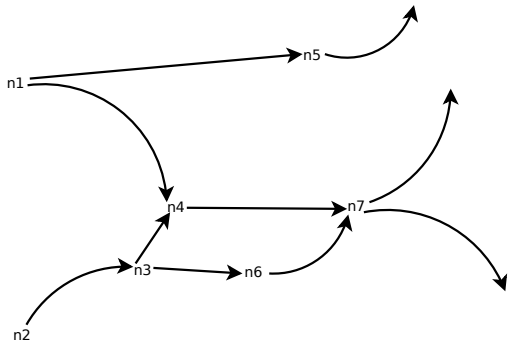

Figure 2: Forklifts on a network

to each platform is a robot arm: t_1 and t_2. It can have different positions relative to its platform, but only within a limited range. Furthermore, we have some kind of internal state for each platform.

The behavior of each robot is controlled by a program. This program implies behavior which is potentially non-deterministic and may depend on other robots or external events. The behavior is characterized by spatial aspects, the movement of the robot, communication aspects (e.g., interactions with other robots or some external controlling device), and internal state changes.

We are interested in simulating possible window cleaning scenarios. Robots have local states and when they interact with another robot their actions may undergo a synchronization. This synchronization does not need to be global, i.e., in case of many robots cleaning a window, the synchronisation does not have to be shared with all robots. For this reason, we only have a partial order of time. Each element of this partial order is called a time point.

One aspect that we are interested in, is whether a certain system state implies a collision. To do this, we could examine the position of each platform and each robot arm. Based on this, we could calculate the exact space boundaries that each device uses. An alternative way is to use an abstraction and use an overapproximation of the space used. More coarse grained abstractions may only have to take the position of the platforms into account.

Forklifts Figure 2 shows a second example. Suppose that two forklifts want to drive on the pictured network of paths and perform tasks. The behavior of the two forklifts is written as follows:

- Forklift 1 starts at n1. It can continue either to n5 or via n4 to n7 in order to exit the system. This is unknown at the time the system starts.
- Forklift 2 starts at n2 and continues to n3. If it can detect that n4 is occupied at will continue to n6 and than n7. If n4 is unoccupied it will continue to n7 via n4. Forklift 2 will exit the system via n7.

Further constraints apply:

- In addition to the movement in the paths network, each forklift occupies a certain amount of space and drives at a certain speed. For distinct time points it is possible to determine an area in which the forklift will reside.
- Additionally each forklift has a locally limited ability to detect obstacles and other forklifts.

Possible questions apply: 1) Can collisions occur? 2) Are forklifts delayed by other forklifts that can be detected and avoided?

Driving Cars Figure 3 shows two cars $c1$ and $c2$ driving on a network of roads. The road network is formalized as a graph. Cars have a limited ability of locally communicating with each other, iff they are within a certain distance of each other. This distance is indicated by a circle around the car in the figure. A car has a local state. This is defined

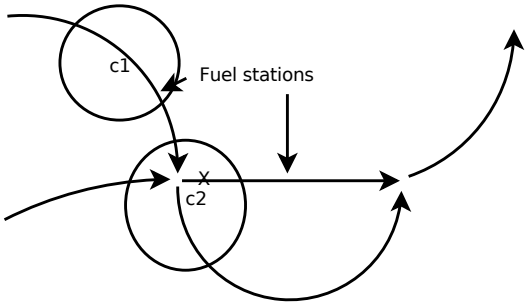

Figure 3: Communicating cars

by the following ingredients:

- The road section it is currently driving on: the edge of the corresponding graph.
- The position on the road, e.g., indicating the progress on the road section and the direction it is facing.
- Aspects of the internal state that encapsulates the parts of a state that are truly local to the car like the amount of fuel left.
- Communication aspects of an internal state, e.g., states that are reached throughout an ongoing communication effort with other cars according to a protocol.

One can see, that spatial, truly internal and communication aspects are encoded in the states.

For calculating the distance between two cars travelling on different roads, we need an interpretation. We take the road edges and the position on the road into account. The topological graph information needs to be interpreted in a geometrical way. For example, we can retrieve data from a geographic information system to get the exact location on

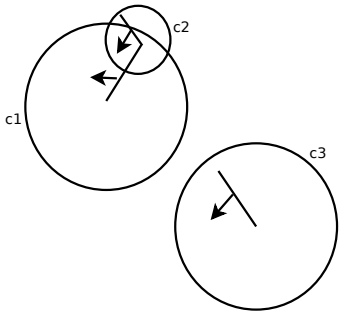

Figure 4: Rotating robots schematic system overview

earth that allows us to compute the distance between the two cars. Alternatively a car may itself keep track of the actual coordinates within its state. Furthermore, we need a shared time points between the two cars so that we have access to their positions at the same time.

The local communication possibilities of the cars may be used for various purposes. For example, a car might inform another car that a road has been blocked (indicated by the X in the figure). The road blockage may leave a fuel station inaccessible. The other car may therefore alter a scheduled fuel stops and use another fuel station.

Rotating Robots Figure 4 depicts an example demonstrating component movement relative to another component. It shows two robots – components $c1$ and $c3$ are performing circular movements. Robot $c1$ features an attachment point. In additional component $c2$: a robot arm is attached to this attachment point. $c2$ also performs circular movements, relative to the movement of $c1$.

A question one wants to investigate is whether a collision between $c3$ and $c1$ is possible.

3.2 System Layers

We propose the use of four different layers for modeling systems and their behavior in terms of space, internal behavior and communication.

1. A topological or geospatial coordinate system. In our window cleaning example this is provided by the window itself. In the communicating cars example this is provided by the topological graph representing the roads and its geographic / geometric interpretation.

2. Static components and their interconnections. In the sense of for example a street network, we can have road blocks and road construction sites as well as refuel stations. In the window cleaning example, we can have single obstacles on the surface.

3. Mobile components. These comprise the cars and robots from our examples moving within or across their allocated segment without changes in the static components structure. Behavior, may be expressed using geometric expressions.

4. Information flow. This can occur between mobile and static components such as cars or robots or a fuel station.

Systems can evolve and change their structure and behavior. A new road may be added or a new fuel station. Cars can move and communicate with each others. The layers are ordered with respect to the rate of expected changes, with layer 1 seeing the least changes and level 4 most changes over time.

3.3 Specification Layers

We allow different layers of abstractions for describing systems and their components. One can mix different specification methods to describe a widely distributed system, e.g., for covering the interactions, spatial behavior (topological and geometric) and aspects of state changes within components.

Specification Formalisms For describing components and their interactions on a requirements level, the following possibilities are reasonable:

- *UML state charts*, *Message Sequence Charts* and *Concurrent automata* can be used to specify the behavior of single components and their interactions.
- *process algebra terms* can be used as another way to specify components and their interactions.
- We can instrument *Pieces of program code* (e.g., a domain specific language, but also plain Scala or Java code can be used). The instrumentation can be used to generate sequences of behavior at runtime and eventually create behavioral descriptions from them.

Furthermore, for describing interactions between components, we use:

- Logical formula describing events and time points implying interactions.
- Partially ordered sets of events or shared time points, are used to describe shared time where different components may interact with each other.

Components As described above, we distinguish static and mobile components, these can be aggregated into subcomponents, e.g., an industrial robot can be made up of arm segments and a mounted tool. Each subcomponent can have its own behavioral descriptions.

Time Time in our models is defined by a set of time points T and a partial order on these time points $\lesssim_T$. Each component specification is associated with a set of time points that is relevant for it. Two time points can form a time interval. Which can serve as a means for a safe overapproximation of behavior that is happening in between two time points.

Properties In particular we are interested in properties that deal with overlapping conditions of space.

- In case of avoiding collisions we are dealing with occupied space for components. The occupied space specified in our abstractions that are the basis for verification and analysis must present an overapproximation of the space that is really used in the real-live system to derive a safe (cf. [27]) result. We require that occupied space of components may overlap at no shared time point or time interval.
- In case of communication possibilities, with limited ranges, we work with space that is a safe underapproximation of a sender and receiver range. Here, we can require space inclusion for a set of shared time points or time intervals.

We can combine safe overapproximations and underapproximations into properties, e.g., to state that a component is always within a certain range of another component, but never comes too close (safety margin) given certain speed limits.

4 Reasoning about Spatial Models

We describe the verification part of our framework in this section.

4.1 Overview

Figure 5 shows the workflow for checking the properties of models with respect to involved computation steps and tools. Models and properties are given to our tool chain for checking.

- In a first step, an algorithm is used to identify large scale verification goals. For example static components that are separated in space and do not interact with each other may be checked independently. The behavior of mobile components may also be checked independently if they are only acting in a local area in space or time. Properties may only regard a local area and therefore only the behavior of components that act within this local area has to be taken into account. Best practices for this approach can be scenario and domain specific. For determining whether one component or a property depends on another, techniques from model checking like cone-of-influence reduction [16] can be applied. A result of this step is the identification of a set of relevant components for the desired properties.
- In a next step, we compute the abstractions of behavior: *invariants* for all relevant components. One way to do this is to unfold all possible execution traces of an

automaton, include all relevant events and annotate them with time points. These traces are than formalized as invariants.

- We use these invariants to generate verification conditions in a next step. Verification conditions can be checked by separate highly specialised tools like SAT [2] and SMT solvers (e.g., Yices [17] or z3 [29]).
- The last step collects the results of used verification tools and presents an overall results. Optionally, we refine invariants or verification conditions, if the result does not satisfy our needs. Eventually this may lead to an iterative process: a fixpoint computation or counterexample guided abstraction refinement [15] where invariants are the predicates.

Computation of invariants can be done in parallel in many cases. Verification conditions never depend on each other, so they can be always checked in parallel.

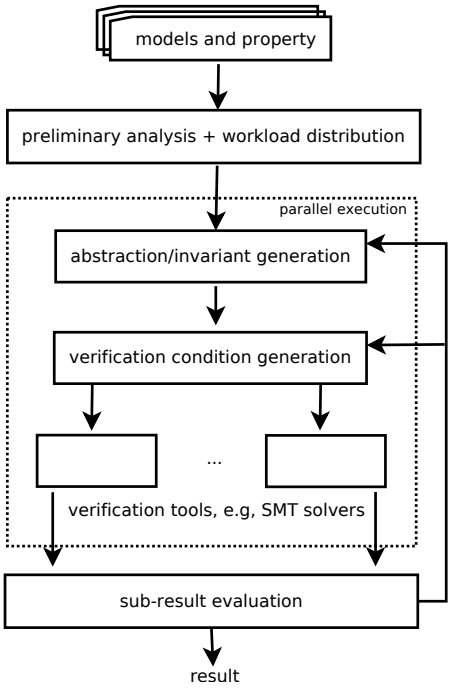

Figure 5: Tool workflow

4.2 Component Oriented Workflow

The checking workflow of possible interference of two components is done in the following way. First, a behavioral abstraction: an invariant is created for each component which is than broken down or combined with information from another component's invariant to

[2] e.g., by using Sat4j: `http://www.sat4j.org/`

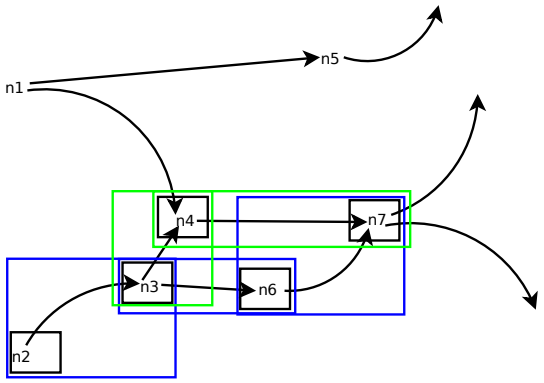

Figure 6: Room occupation invariants (time points)

derive verification conditions. These are given to a solver. Here, invariants contain only overapproximations of single components and do not take restrictions due to possible interference with other components into account. Thus, overapproximations are safe, but may be coarse.

For achieving more fine grained approximations, optionally, we can create logical conditions capturing the interference of components. These are invariants, too. Based on them, we can create verification conditions that further restrain verification conditions derived solely from the invariants that capture the behavior of single components.

4.3 Invariants

Our invariants capture the entire behavior of a component over time or distinct aspects of interference between components. They provide a safe overapproximation of behavior.

Nature of Invariants Like invariants in classical Hoare style loop verification, our *spatial component invariants* are supposed to hold over the entire execution time of a system. Like loop invariants, spatial component invariants typically abstract from the actual implementation. Similar to loop invariants, where variables' values are often described in relation to a changing loop index, we describe the occupied or otherwise classified space and if applicable other component state based information in relation to time.

Invariants for Occupied Space Figure 6 shows invariants for the second forklift in our forklift example. These overapproximate the occupation of space for time points that are associated with the nodes. Furthermore, overapproximations for time intervals between time points are shown. In this case the invariants provide a geometric interpretation of the underlying graph. Here, to gain, a safe collision analysis, we must use overapproximations of the space that is actually occupied.

Invariants for Communication and Detection Figure 7 shows invariants for the second forklift in the forklift example indicating visibility ranges to detect other forklifts. In

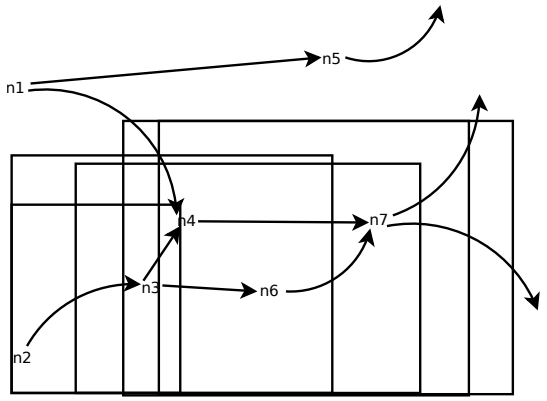

Figure 7: Room occupation invariants (time intervals)

order to safely analyse detection possibilities, we need to work with underapproximations. Invariants characterising time intervals are not shown, but are created from union operations between the invariants that characterize two neighbouring time points.

Note, that for displaying purposes, the granularity of time points in Figure 6 and Figure 7 is large. For real systems, the boxes are much closer together, so that safe overapproximations can be much tighter to the actual system behavior.

Invariants for Interplay between Components On an abstract view, we can associate interaction possibilities between components with events. Events can also be associated with detection of obstacles or external interactions. These events can be used in invariants, to characterize interactions between components. Such a formula can have the form:

$$(\text{ time point } \wedge \text{ event }) \longrightarrow \text{occupied space \& other actions}$$

We can create invariants that describe aspects of interactions between components. Some possible conditions on triggering events can have the forms

$$\text{time point} \rightarrow \text{event}$$
$$\text{or}$$
$$(\text{ time point } \wedge \text{ event }) \rightarrow \text{event'}$$

Combining these invariants can be done by intersection with other invariants characterizing the behavior of a single component in order to restrain its behavior and derive more refined verification conditions that are given to solvers.

Aggregating Subcomponents We can aggregate subcomponents into larger components. Likewise, we can integrate their invariants. Figure 8 shows invariants characterizing occupied space for a given point in time for the cleaning platforms from Figure 1. The right side displays invariants for subcomponents like the actual platform, the robot arm and the tool while the left side displays a safe overapproximation of the combined system.

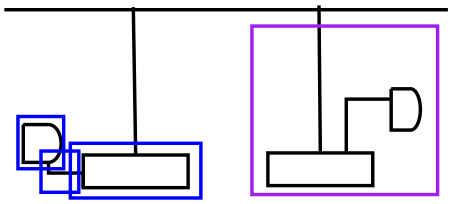

Figure 8: Concurrent window cleaning

Term based Representations of Invariants We represent invariants syntactically as terms. Invariants are built from predicates that are combined using function symbols like logical operators. An invariant that characerizes the spatial behavior of a single component has the form:

$$\begin{array}{l} \ldots \\ t = i \quad \rightarrow \mathit{OccupySpace}_i \quad \wedge \mathit{CommunicationRange}_i \quad \wedge \ldots \\ t = i+1 \rightarrow \mathit{OccupySpace}_{i+1} \wedge \mathit{CommunicationRange}_{i+1} \wedge \ldots \\ t = i+2 \rightarrow \mathit{OccupySpace}_{i+2} \wedge \mathit{CommunicationRange}_{i+2} \wedge \ldots \\ \ldots \end{array}$$

For non linear time, the i index is replaced by a more complex time point indicating the node in the graph characterizing the partial order of time points. We can perform operations on invariants like the merging of time points into time intervals thereby transforming it into another syntactic representation. For the example above this is realized as follows:

$$\begin{array}{l} \ldots \\ t \in [i, i+1) \rightarrow \\ \qquad \mathit{OccupySpace}_i \sqcup \mathit{OccupySpace}_{i+1} \wedge \\ \mathit{CommunicationRange}_i \sqcup \mathit{CommunicationRange}_{i+1} \wedge \ldots \\ t \in [i+1, i+2) \rightarrow \\ \qquad \mathit{OccupySpace}_{i+1} \sqcup \mathit{OccupySpace}_{i+2} \wedge \\ \mathit{CommunicationRange}_{i+1} \sqcup \mathit{CommunicationRange}_{i+2} \wedge \ldots \\ \ldots \end{array}$$

Here, $\rightarrow$ and $\wedge$ are function symbols of the term. The $\sqcup$ is a function that performs an interpretation and yields a subterm representing a semantic union of its arguments.

Furthermore, we have operations that comprise simplifications, abstractions, normalizations and assigning and removing ownership – means to indicate that space belongs to a distinct component – to occupied space. Contrary to abstractions we have interpretations on the invariant level, e.g., for assigning a geometry to parts of topological or high-level component representations, thereby lifting invariants to a more concrete level.

Generating Invariants Invariants may be directly encoded or generated in the following ways from spatial behavioral specifications.

- *Unfolding of automata to generate timed conditions in invariants.* If we are only interested in checking a limited time span, this is straight forward. For periodically repeating behavior, we can group time points into equivalence classes for a finite invariant that describes the infinite behavior. A special case of automata unfolding are sequences of states that a system component has gone through during a simulation. These sequences can be generated from other tools such [3].

- *Instrumentation of code pieces to generate timed conditions in invariants at runtime during a testing phase of a system.* Here, we have to ensure that we cover all relevant aspects of a system, because we are in danger of getting an unsafe underapproximation of the behavior. However, regarding the behavior of robots, it seems reasonable to exhaustively catch all spatial movements of a robot without having exhausted the robot's software state space.

- *Transformation of specifications on a symbolical / semantical level.* This involves symbolic analysis of the program code that controls a distributed system and the extraction of invariants based on them.

Abstractions using Invariants We can use abstraction on the invariant level to facilitate later verification tasks. For space, we can use boxes to over- or underapproximate more detailed structures. Figure 9 shows a screenshot of a box based overapproximation in the rotating robots example for collision detection. Here, the exact amount of space for segments of robot components is automatically abstracted using overapproximations.

4.4 Verification Conditions

We connect different tools including SAT and SMT solvers to our framework. Verification conditions are generated from invariants so that a property can be verified and the input is suitable for the particular tool.

Predicate Based Space Conditions In our invariants we can use predicates to characterize the occupation of space for a distinct time point or time interval. As an example: each point indicating the occupation of a square on a plane can be identified using a predicate. By doing so, space occupied by different components can be made comparable on the invariant level. Shown in Figure 10 is the unfolding of a predicate characterizing the occupation of a box specified by an upper left and a lower right pair of coordinates. The resulting representation can be compared with unfolded invariants of other components. This comparison must be made for each shared time point or time interval between components and can by either given to a SAT solvers or can by realized by a small algorithm implemented in Scala in our framework.

[3] We have successfully applied it together with Reactive Blocks and its specification method [24]

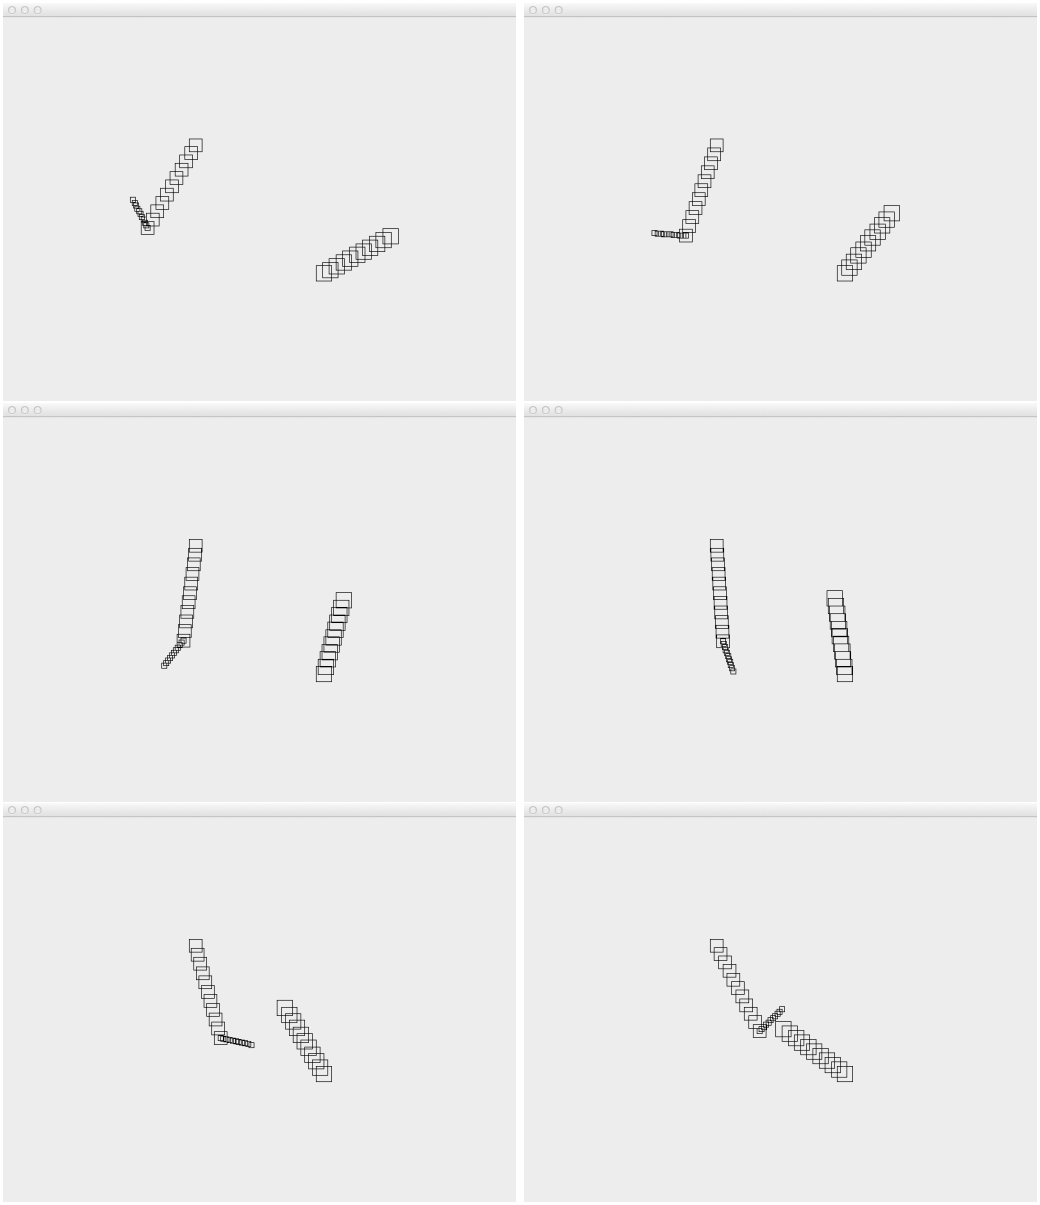

Figure 9: Box-based overapproximation (Screenshots)

$$
\begin{aligned}
&occupyXYspace(x_1, y_1, x_2, y_2) = \\
&\quad occupyXY(x_1, y_1) \wedge occupyXY(x_1 + 0.1, y_1) \wedge ... \wedge occupyXY(x_2, y_1) \wedge \\
&\quad occupyXY(x_1, y_1 + 0.1) \wedge occupyXY(x_1 + 0.1, y_1 + 0.1) \wedge ... \wedge occupyXY(x_2, y_1 + 0.1) \wedge \\
&\qquad ... \\
&\quad occupyXY(x_1, y_2 - 0.1) \wedge occupyXY(x_1 + 0.1, y_2 - 0.1) \wedge ... \wedge occupyXY(x_2, y_2 - 0.1) \wedge \\
&\quad occupyXY(x_1, y_2) \wedge occupyXY(x_1 + 0.1, y_2) \wedge ... \wedge occupyXY(x_2, y_2)
\end{aligned}
$$

Figure 10: Predicate unfolding (conjunction of points)

Space Conditions Based on Inequalities A second way of comparing different invariants that characterize geometric occupation of space is to generate for each shared time point a geometric formula characterizing the potential overlapping of occupied space. For example, two components c_1 and c_2 specified with a spatial occupation specified by boxes with upper left and lower right coordinates as given below:

$$(x_{1_c}(t), y_{1_c}(t)),\ (x_{2_c}(t), y_{1_c}(t)),$$
$$(x_{1_c}(t), y_{2_c}(t)),\ (x_{2_c}(t), y_{2_c}(t))$$

Where each component is given by a function that depends on a time parameter t we generate the following:

$$\neg\exists x, y.\quad \bigvee_{t \in \textit{shared time points}} (x_{1_{c_1}}(t) \leq x \leq x_{2_{c_1}}(t) \wedge x_{1_{c_2}}(t) \leq x \leq x_{2_{c_2}}(t) \wedge y_{1_{c_2}}(t) \leq y \leq y_{2_{c_2}}(t) \wedge y_{1_{c_2}}(t) \leq y \leq y_{2_{c_2}}(t))$$

A solution characterizes the overlapping at any shared time point. If no solution exists, there is no overlapping. We can use this as a verification condition thereby unfolding the $\bigvee_{t \in \textit{shared timepoints}}$ and treating multiple time points or generate an independent verification condition for each time point.

5 Implementation and Application

Here, we describe some features of our framework that we have implemented and evaluated.

5.1 Specification of Models and Properties

Invariants and operations on them are encoded using abstract data type style *case classes* from the Scala language.

Basic Constructs To give a look and feel, a small excerpt of the abstract datatypes in Scala is given below:

```
abstract class Invariant;

abstract class ATOM extends Invariant;

case class OR (t1 : Invariant, t2 : Invariant)
   extends Invariant;
case class AND (t1 : Invariant, t2 : Invariant)
   extends Invariant;
case class NOT (t : Invariant) extends Invariant;
case class IMPLIES (t1 : Invariant, t2 : Invariant)
   extends Invariant;
        ...
```

```
case class BIGOR (t : List[Invariant]) extends Invariant;
case class BIGAND (t : List[Invariant]) extends Invariant;
        ...
case class TimePoint [T](timepoint : T)
   extends ATOM;
case class TimeInterval [T]
 (timepoint1 : T, timepoint2 : T)
   extends ATOM;
case class Event[E] (event : E) extends ATOM;
        ...
case class Occupy3DBox
  (x1 : Int, y1: Int, z1 : Int,
   x2 : Int, y2 : Int, z2 : Int) extends ATOM;
case class OccupySegment3D
  (x1 : Int, y1 : Int, z1 : Int,
   x2 : Int, y2 :Int, z2 : Int, radius : Int)
     extends ATOM;
case class Occupy3DPoint (x:Int, y:Int, z: Int)
     extends ATOM
```

The excerpt shows only some logical operations, time (the actual type used for representing time is parameterized `T`, any type with a partial order is suitable) and geometric constructs for an Euclidian space. Some abbreviations for different levels of modeling and easier automatic processing can be seen. For example, the `BIGOR` and `BIGAND` constructs are semantically equivalent to nested `AND` and `OR` constructs.

Ownership of Geometric Structures We support the ownership of geometric structures to identify structures that belong to a specific aspect of an entity. For example, the following constructors are available:

```
case class OwnBox[C] (owningcomponent : C,x1 : Int,y1 : Int,x2 :
Int,y2 : Int) extends ATOM;
// a box that is owned by a component.
        ...
case class OwnPoint[C] (owningcomponent : C,x:Int, y:Int) extends ATOM
```

Relative Time Points and Symbolic Integer Values In addition to absolute time, we also allow the specification of time relative to an event (event relative time point `ERTP`) and the use of symbolic values as integers for coordinates (type `SI`. Geometric constructs making use of these types are provided. For example, the following constructors are available:

```
abstract class TIMEPOINT;

abstract class ERTP; // Event Relative Time Point

case class TERTP [E,T] (event : E, offset : T) extends ERTP;
case class IntERTP [E](event : E, offset : Int) extends ERTP;

abstract class SI; // SymbolicInt
case class SI_C (c:Int) extends SI;
case class SI_Add (i1 : SI, i2 : SI) extends SI;
case class SI_Sub (i1 : SI, i2 : SI) extends SI;
```

```
case class SI_Times (i1 : SI, i2 :SI) extends SI;
case class SI_Var [V](v : V) extends SI
        ...
case class OccupyBoxSI (x1 : SI,y1 : SI,x2 : SI,y2 : SI) extends ATOM;
// symbolic coordinates, alternative concept to EROccupyBox
case class EROccupyBox (
   x1 : (ERTP => Int),y1 : (ERTP => Int),
   x2 : (ERTP => Int),y2 : (ERTP => Int)) extends ATOM;
// Coordinates depending on an event relative time point
```

Specification Example A very small example for an invariant encoded in Scala describing only the possible movement of a forklift in the graph shown in Figure 2 is presented below (here, some more abstract representation `OccupyNode` is chosen):

```
  val topologicalinvariant_fl2 : Invariant =
   AND( AND ( AND (
      IMPLIES(TimePoint("pt1"),OccupyNode("n2")),
      IMPLIES(TimePoint("pt2"),
         OR (OccupyNode("n3"),OccupyNode("n4")))),
      IMPLIES(TimePoint("pt3"),
         OR (OccupyNode("n6"),OccupyNode("n7")))),
      IMPLIES(TimePoint("pt4"),OccupyNode("n7")));
```

Only 4 different time points and no geometric interpretation are encoded.

In realistic examples, we use functions that build abstract datatypes comprising thousands of constructors based. The construction can be based on parameters to characterize different behaviors and interaction possibilities of different components.

5.2 Parameterized Specifications

Parameterized specifications can be created by Scala functions that create distinct formalizations based on their parameters.

Lifting of a Workpiece Figure 11 shows a grapple hook lifting a workpiece. The behavior of the hook and the workpiece is formalized in BeSL. Here, a grapple hook consists of four subcomponents: the two arms for holding the workpiece on the sides, a components that links the two arms and a chain that can be lowered or raised to position the hook. The workpiece is just formalized as a single component.

Parameterized Lifting Invariants We have experimented with different behavioral descriptions for the hook. One parameterized description for one of the arms is given below. It formalizes the behavior of for 200 time points. The speed can be given as a parameter. Furthermore, the lifting may be stopped for a given point in time.

```
def invariantR3(speed : Float, stoppointup: Int) : Invariant ={
var inv : List[Invariant] = Nil;

for (i <- 0 to 100) {
  inv ::= (IMPLIES(TimeStamp (i),
      OccupySegment(300,200+(i*speed).toInt,320,200+(i*speed).toInt,3)))
}
```

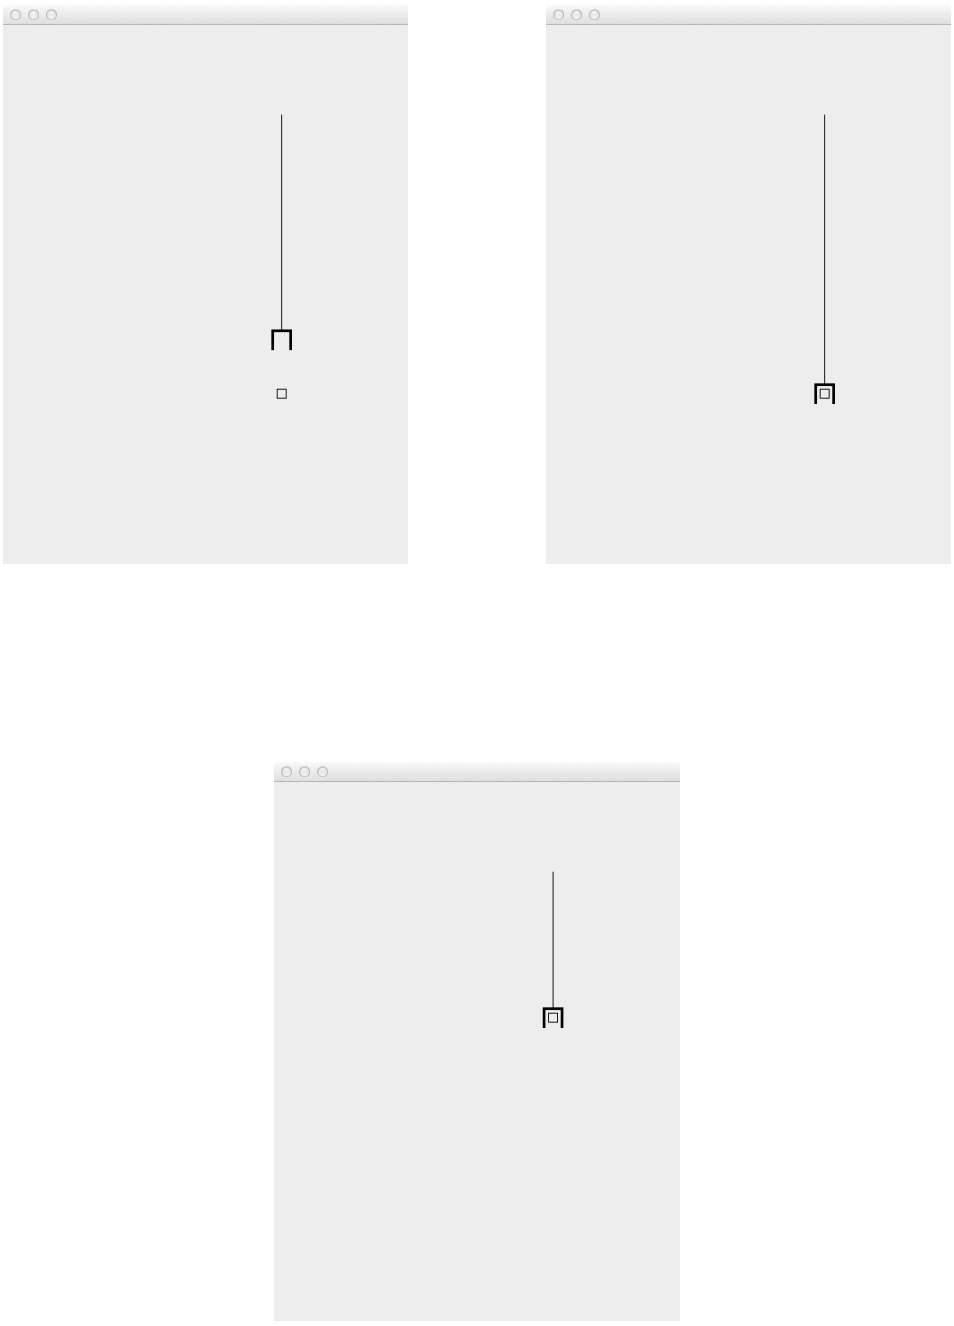

Figure 11: Workpiece lifting

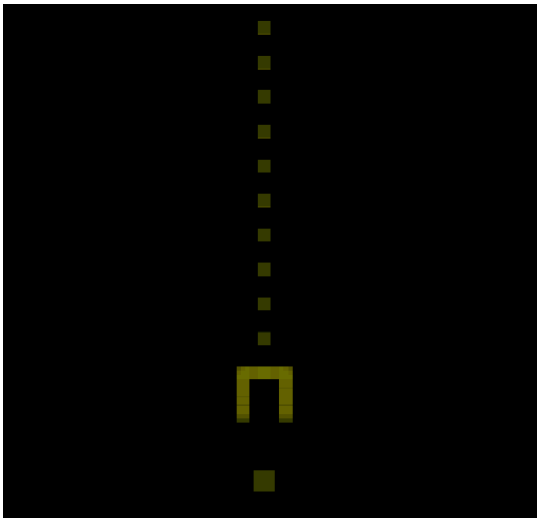

Figure 12: Workpiece lifting 3D

```
for (i <- 0 to 100) {
     if (i < stoppointup) {
     inv ::= (IMPLIES(TimeStamp (i+101),
     OccupySegment(
                     300,
                     200+(100*speed).toInt-(i*speed).toInt,
                     320,
                     200+(100*speed).toInt-(i*speed).toInt,3)))
     } else {
     inv ::= (IMPLIES(TimeStamp (i+101),
     OccupySegment(
                     300,
                     200+(100*speed).toInt-(stoppointup*speed).toInt,
                     320,
                     200+(100*speed).toInt-(stoppointup*speed).toInt,3)))
     }
}
return (BIGAND(inv));
}
```

A 3D version of the workpiece lifting is shown in Figure 12. It contains an abstraction of segment descriptions based on 3D boxes.

5.3 Checking Collision Conditions

We have implemented the checking of component invariants for possible collisions. As a more realistic example and means for measuring performance we have regarded two invariants for collisions where each invariant comprises 1000 different time point entries

and associated geometric occupied space.[4]

1. Breaking the invariants down to SMT verification conditions where one verification condition is generated per time point and calling the SMT solver takes between 20 and 25 seconds.

2. Breaking the invariants down to a single SMT verification condition and checking can be performed in slightly less then one second.

3. Breaking the invariants down to verification conditions each one comprising 15000 predicates indicating the occupation of single points for one invariant and 20000 predicates indicating the occupation of single points for the second invariant for each time point can be done in at most 7-8 seconds. Here, a HashSet based Java implementation is used. In case of early collision detection the checking time can be significantly less.

From experiments 1 and 2 one can be seen that it is usually a good idea to encode as many information as possible into an SMT formula, since the overhead of calling it is significant.

If higher resolutions are used, the single points' abstractions in the third experiment performance decreases linear with the amount of points used. The third experiment can be done using a SAT solver, this takes significantly longer, due to the higher complexity of SAT formula compared to the Java HashSet and the less abstraction power compared to the SMT based approach.

Collisions have to be checked between each two components for each shared time point or time interval. Our notion of time allows the characterization of components that will never interfere, so the amount of components that one has to check for a time point does depend on the application scenario. Analysing whether a time point or time interval is shared or not can be based on requirements and specification documents. To achieve a safe result, we must use an overapproximation of shared time points and time intervals.

Furthermore, it is possible that components have a non-deterministic behavior. In that case, we need to adapt out SMT conditions for the experiments one and two or use the approach based on the experiment three and add all possible occupied point to the HashSet.

Visualization In our framework single invariants and scenarios consisting of multiple invariants can be graphically depicted as demonstrated in the figures throughout the report. In Figure 13 the x and y coordinates of three objects' spaces are plotted against the time.

[4]The following checking times where retrieved for an intel core i5 running 2.8 GHz with 8 GB RAM running Mac OS 10.8.4. The SMT solver used is z3 [29].

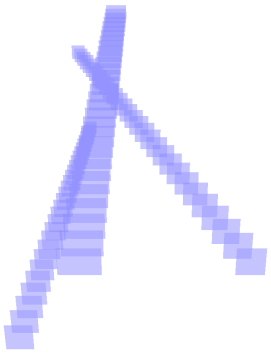

Figure 13: 3D graphical invariant representation

6 Conclusion and Future Work

We presented work towards a unified approach to cyber-physical systems and component-based software development by motivating a new framework concentrating on spatial aspects of cyber-physical systems. We described a process for checking properties of our models and described the approach using different examples.

Different directions for future work and ongoing work comprise: 1) Taking the parallel checking of our models into account and implementing it for a suitable platform. 2) Offering SOA based services for spatial modeling and checking. 3) Building a spatial behavioral type system for components following our previous ideas [7, 6].

Another more academic area of future work is reasoning about spatial properties in proof assistants like Coq[5] or Isabelle [30]. Here, we are aiming at 1) connecting proof assistants as additional reasoning tools to manually verify concrete properties at an invariant or verification condition level that are undecidable 2) derive meta-results on the process of getting from models to invariants and verification conditions.

References

[1] M. Aiello, I. E. Pratt-Hartmann, and J. FAK van Benthem, eds. Handbook of spatial logics. Springer, 2007.

[2] M. Apel. A 3D geological information system framework. Geophysical Research Abstracts, vol. 7, European Geosciences Union, 2005.

[3] A. Basu, M. Bozga, J. Sifakis. Modeling Heterogeneous Real-time Components in BIP. Fourth IEEE International Conference on Software Engineering and Formal Methods, IEEE 2006.

[4] B. Bennett, A. G. Cohn, F. Wolter, M. Zakharyaschev. Multi-Dimensional Modal

[5] `http://coq.inria.fr/`

Logic as a Framework for Spatio-Temporal Reasoning. Applied Intelligence, Volume 17, Issue 3, Kluwer Academic Publishers, November 2002.

[5] S. Bensalem, A. Griesmayer, A. Legay, T. H. Nguyen, J. Sifakis, R. Yan. D-finder 2: towards efficient correctness of incremental design. In NASA Formal Methods, LNCS, Springer, 2011.

[6] J. O. Blech. Towards a Framework for Behavioral Specifications of OSGi Components. Formal Engineering approaches to Software Components and Architectures. Electronic Proceedings in Theoretical Computer Science, 2013.

[7] J. O. Blech, Y. Falcone, H. Rueß, B. Schätz. Behavioral Specification based Runtime Monitors for OSGi Services. Leveraging Applicationsof Formal Methods, Verification and Validation (ISoLA), vol. 7609 of LNCS, Springer, 2012.

[8] J. O. Blech and M. Périn. Generating Invariant-based Certificates for Embedded Systems. ACM Transactions on Embedded Computing Systems (TECS), 2012.

[9] J. O. Blech and H. Schmidt. Towards Modeling and Checking the Spatial and Interaction Behavior of Widely Distributed Systems. Improving Systems and Software Engineering Conference, Melbourne, 2013.

[10] G. Booch, J. Rumbaugh, and I. Jacobson. The Unified Modeling Language User Guide (second edition). Addison-Wesley, 2005.

[11] L. Caires and L. Cardelli.A Spatial Logic for Concurrency (Part I). Information and Computation, Vol 186/2 November 2003.

[12] L. Caires and L. Cardelli. A Spatial Logic for Concurrency (Part II). Theoretical Computer Science, 322(3) pp. 517-565, September 2004.

[13] L. Caires and H. Torres Vieira. SLMC: a tool for model checking concurrent systems against dynamical spatial logic specifications. Tools and Algorithms for the Construction and Analysis of Systems. Springer, 2012.

[14] J. Canny. The complexity of robot motion planning. MIT Press, Cambridge, 1988.

[15] E. Clarke, O. Grumberg, S. Jha, Y. Lu, H. Veith. Counterexample-Guided Abstraction Refinement. Computer Aided Verification, vol. 1855 of LNCS, Springer, 2000.

[16] E. Clarke, O. Grumberg, D. A. Peled. Model Checking. MIT Press, 1999.

[17] B. Dutertre, L. De Moura. The yices smt solver. Tool paper at http://yices.csl.sri.com/tool-paper.pdf, 2006.

[18] F. Gadducci and A. Lluch Lafuente. Graphical Verification of a Spatial Logic for the π-calculus." Electronic Notes in Theoretical Computer Science 154.2, 2006.

[19] S. Gottschalk , M. C. Lin , D. Manocha , S. Gottschalk , M. C. Lint , D. Manocha. OBB-Tree: A hierarchical structure for rapid interference detection. Proc. ACM SIGGRAPH, 171-180, 1996.

[20] F. Han, J. O. Blech, P. Herrmann, H. Schmidt. Towards Verifying Safety Properties of Real-Time Probabilistic Systems. Formal Engineering approaches to Software Components and Architectures, 2014. *to appear*

[21] D. Hirschkoff, É. Lozes, D. Sangiorgi. Minimality Results for the Spatial Logics. Foundations of Software Technology and Theoretical Computer Science, vol 2914 of LNCS, Springer, 2003.

[22] P. Jiménez, F. Thomas, C. Torras. Collision Detection Algorithms for Motion Planning. Lectures Notes in Control and Information Sciences 229, Springer, 1998.

[23] S. Kambhampati and L.S. Davis. Multiresolution path planning for mobile robots. Volume 2 , Issue: 3, Journal of Robotics and Automation, IEEE 1986.

[24] F. A. Kraemer and P. Herrmann. Service Specification by Composition of Collaborations–An Example. International conference on Web Intelligence and Intelligent Agent Technology. IEEE, 2006.

[25] J-C. Latombe. Robot Motion Planning. Kluwer Academic Publishers, 1991.

[26] M. C. Lin, D. Manocha. Fast interference detection between geometric models. The Visual Computer, Volume 11, Issue 10, pp 542-561, Springer, 1995.

[27] C. Loiseaux, S. Graf, J. Sifakis, A. Bouajjani and S. Bensalem. Property preserving abstractions for the verification of concurrent systems. Formal Methods in System Design, Volume 6 Issue 1, Kluwer Academic Publishers, 1995.

[28] G. Frehse, C. Le Guernic, A. Donzé, S. Cotton, R. Ray, O. Lebeltel, R. Ripado, A. Girard, T. Dang, O. Maler. SpaceEx: Scalable Verification of Hybrid Systems. Computer aided verification (CAV'11), 2011.

[29] L. De Moura, N. Bjørner. Z3: An efficient SMT solver. In Tools and Algorithms for the Construction and Analysis of Systems (pp. 337-340). Springer, 2008.

[30] T. Nipkow, L. C. Paulson, M. Wenzel. A Proof Assistant for Higher-Order Logic. Vol. 2283 of LNCS, Springer, 2002.

[31] A. Platzer, J-D. Quesel. KeYmaera: A Hybrid Theorem Prover for Hybrid Systems (System Description). International Joint Conference on Automated Reasoning, vol. 5195 of LNCS, pages 171-178. Springer, 2008.

[32] R. H. Reussner, I. H. Poernomo, H. W. Schmidt. Reasoning about Software Architectures with Contractually Specified Components. Component-Based Software Quality, vol. 2693 of LNCS, Springer, 2003.

[33] G. Smith and J. Friedman. A Technology Whose Time Has Come. Earth Observation Magazine, November 2004.

[34] A. Troynikov and H. Schmidt. Designing robust fault tolerant systems with shape. Improving Systems and Software Engineering Conference incorperating SEPG^{SM} Asia-Pacific Conference 2012, Melbourne, August 2012.

[35] C. Weaver, D. Peuquet, A. M. MacEachren. STNexus: An Integrated Database and Visualization Environment for Space-Time Information Exploitation. `http://www.geovista.psu.edu/publications/2005/Weaver_ARDA_05.pdf`, 2005.

[36] S. Dal Zilio, D. Lugiez, C. Meyssonnier. A logic you can count on. Symposium on Principles of programming languages, ACM, 2004.